

TRAN-P082

UNITED STATES PATENT APPLICATION

for

EXPLICIT CONTROL OF SPECULATION

Inventors:  
H. PETER ANVIN  
DAVID DUNN

Prepared by:

WAGNER, MURABITO & HAO LLP  
TWO NORTH MARKET STREET  
THIRD FLOOR  
SAN JOSE, CALIFORNIA 95113  
(408) 938-9060

## EXPLICIT CONTROL OF SPECULATION

### BACKGROUND OF THE INVENTION

#### FIELD OF THE INVENTION

- 5           Embodiments of the present invention relate to processors and microprocessors. More specifically, embodiments of the present invention relate to speculative execution of operations.

#### RELATED ART

- 10           Speculative execution is known in the art of microprocessors. Should an event occur where speculation is not permitted, speculation is suspended while the event is handled. A fault provides one example of such an event.

- When implementing a complex instruction set computer (CISC) instruction  
15 set, some operations require the execution of microcode of some sort. Generally, it is desirable for the microcode to have access to the speculation function. However, when speculation is suspended, this will not be the case.

## SUMMARY OF THE INVENTION

Embodiments of the present invention provide methods and systems that allow partial speculation (e.g., speculation within constraints) in situations where speculation is not conventionally permitted.

5

## BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of these embodiments.

5

Figure 1 is a block diagram of an example of a computer system architecture that can be used with embodiments of the present invention.

Figure 2 is a block diagram of another example of a computer system architecture that can be used with embodiments of the present invention.

10

Figure 3 illustrates the use of controlled (partial) speculation according to an embodiment of the present invention.

Figure 4 is a flowchart of a method providing partial speculation according to an embodiment of the present invention.

15

Figure 5 is a flowchart of another method providing partial speculation according to an embodiment of the present invention.

20

## DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the various embodiments of the invention, examples of which are illustrated in the accompanying drawings. While the invention will be described in conjunction with these embodiments, it  
5 will be understood that they are not intended to limit the invention to these embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents, which may be included within the spirit and scope of the invention as defined by the appended claims. Furthermore, in the following detailed description of the present invention, numerous specific details  
10 are set forth in order to provide a thorough understanding of the present invention. However, it will be recognized by one of ordinary skill in the art that the present invention may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail as not to unnecessarily obscure aspects of the present  
15 invention.

Some portions of the detailed descriptions that follow are presented in terms of procedures, logic blocks, processing, and other symbolic representations of operations on data bits within a computer memory. These  
20 descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, logic block, process, etc., is here, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps are those requiring physical manipulations of physical  
25 quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven

convenient at times, principally for reasons of common usage, to refer to these signals as bits, bytes, values, elements, symbols, characters, terms, numbers, or the like.

5           It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "operating,"  
10 "exiting," "entering," "permitting," "suspending," "returning," "counting," "experiencing," "rolling back," "detecting," "handling" or the like, refer to the action and processes (e.g., flowcharts 400 and 500 of Figures 4 and 5, respectively) of a computer system or similar intelligent electronic computing device, that manipulates and transforms data represented as physical  
15 (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

20           Aspects of the present invention may be practiced on a computer system that includes, in general, a central processing unit (CPU) for processing information and instructions, random access (volatile) memory (RAM) for storing information and instructions, read-only (non-volatile) memory (ROM) for storing static information and instructions, a data storage device such as a magnetic or  
25 optical disk and disk drive for storing information and instructions, an optional user output device such as a display device (e.g., a monitor) for displaying information to the computer user, an optional user input device including

alphanumeric and function keys (e.g., a keyboard) for communicating information and command selections to the processor, and an optional user input device such as a cursor control device (e.g., a mouse) for communicating user input information and command selections to the processor. The computer system  
5 may also include a device for providing a physical communication link between the computer system and a network, using either a wired or a wireless communication interface.

Figure 1 is a block diagram of an example of a computer system 100 that  
10 can be used with embodiments of the present invention. In the present embodiment, computer system 100 includes a microprocessor 103 coupled to main memory 101 and to an input/output (I/O) device 106. Main memory 101 functions as the primary workspace of a computer system and can include the operating system software and application software used by the computer  
15 system. I/O device 106 can be any of a variety of peripheral devices other than main memory.

Figure 2 is a block diagram of another example of a computer system 200 that can be used with embodiments of the present invention. Relative to the  
20 architecture of Figure 1, computer system 200 introduces a bus interface 202 and a north bridge 204 between main memory 101, I/O device 106 and microprocessor 103.

Referring to both Figures 1 and 2, microprocessor 103 is generally  
25 described as a complex instruction set computer (CISC) microprocessor. Microprocessor 103 includes CPU core 105, which includes a CPU register 102. CPU register 102 represents an instance of microprocessor memory that can be

referred to as "private memory." As used herein, private memory refers to memory that is not subject to direct memory access (DMA) and that is hidden from input/output (I/O) bus(es). Private memory also refers to memory that is not observable outside of the microprocessor 103, either physically or logically.

- 5 There can also be instances of private memory other than or in addition to CPU register 102.

It is appreciated that computer systems 100 and 200 can include additional components as described above, and that there can be paths between  
10 components different from or in addition to the paths illustrated in Figures 1 and 2.

An I/O interface boundary is represented in Figures 1 and 2. The position of the I/O interface boundary can be different than that shown. The I/O interface  
15 boundary is used here as an abstraction, to illustrate that a memory operation that results in the I/O interface boundary being crossed is externally visible (relative to microprocessor 103) while a memory operation that does not result in the I/O interface boundary being crossed is not externally visible (relative to microprocessor 103).

20

For the purpose of explanation, the discussion herein refers to an "abstract machine." In this context, an abstract machine is a theoretical construct implementing the formal definition of an instruction set architecture (ISA) as visible to operating system and application software, for example a variant of the  
25 x86 (80x86) instruction set architecture.

In one embodiment, microprocessor 103 is an x86 processor (e.g., a microprocessor that implements the x86 ISA using a combination of hardware circuits and microcode). In another embodiment, microprocessor 103 is a processor that transfers to software those functions that are keyed to determining  
5 what instructions (micro-instructions) to execute and when. In one example of the latter embodiment, conventional x86 (80x86) instructions are converted into microarchitecture instructions, for example Very Long Instruction Word (VLIW) instructions using "code morphing™ software."

10 One implementation of code morphing™ software utilizes an interpreter module and a translator module. Among other functions, the interpreter module interprets ISA (e.g., x86) instructions one at a time, similar to that of conventional processors. The translator module is invoked when critical and/or frequently used ISA instruction sequences are detected. In general, the translator module  
15 recompiles the ISA instructions into native instructions that can reduce the number of instructions executed and that can schedule the instructions to execute more efficiently within the microprocessor.

As used herein, a "translation" refers to a microarchitecture instruction or  
20 sequence of microarchitecture instructions that perform the same function as some set of ISA-specified instructions. A translation can be viewed as having a beginning state and an ending state (the ending state is, in essence, the beginning state of the next translation). The beginning and ending states may be referred to as "commit points." At a commit point, the "architecturally visible"  
25 resources of the translating processor and of the abstract machine should be the same.

Embodiments of the present invention introduce methods and systems thereof that allow partial speculation (e.g., speculation within constraints). With partial speculation, after certain types of events are detected, speculation remains enabled for CPU registers and other memories private to a microprocessor, while speculation normally permitted for certain other operations is suspended. Accordingly, while the event is dispatched, some speculation is permitted as opposed to suspending all speculation.

Figure 3 illustrates the use of controlled (e.g., partial) speculation according to an embodiment of the present invention. Execution begins at a point identified as a "speculation boundary." A speculation boundary refers to a memory state of the microprocessor (e.g., microprocessor 103 of Figures 1 and 2) that is present at a particular point in time.

In the example of Figure 3, in step 1, execution proceeds forward from the speculation boundary 310 in a first mode referred to as full speculation mode. Operations can be roughly categorized into three categories: register, memory, and I/O. In full speculation mode, register operations can be speculated as long as their original state can be recovered. Memory operations can be speculated, but in the event of a DMA request, speculation is constrained. Speculation of I/O operations is permitted in some instances and constrained in others.

Specifically, in one embodiment of full speculation mode, speculation is permitted for register operations, operations that involve memory that is private to the microprocessor 103 of Figures 1 and 2 (e.g., memory hidden from an I/O bus or memory protected against DMA), I/O writes but not I/O reads, and main memory reads and writes. Generally, speculation is not permitted in the event of

a DMA request, although there may be exceptions. Also, speculation permits the presence of non-architectural faults (exceptional conditions that do not correspond to observable events in the abstract machine mentioned above).

5        In step 2 of Figure 3, some type of event is identified. An event can be a fault such as an attempt to access a memory page that does not exist, or an attempt to access a segment descriptor that is not valid. An event can instead be a DMA request or an I/O read. In the presence of these kinds of events, full speculation mode may interfere with the externally visible behavior of  
10    microprocessor 103 (Figures 1 and 2), and so for some events such as those mentioned above, full speculation mode should be suppressed.

      In step 3 of Figure 3, roll back to the speculation boundary 310 occurs, restoring the memory state that existed at that point.

15

      In step 4 of Figure 3, execution proceeds forward from the speculation boundary 310 in a second mode referred to as partial speculation mode. In one particular implementation, the second mode is known as the X PROMOTE mode.

20        In partial speculation mode, speculation is permitted for a non-null subset of the operations permitted in full speculation mode. Specifically, in one embodiment of partial speculation mode, speculation is permitted for register operations and for operations that involve memory that is private to the microprocessor 103 of Figure 1 and 2 (e.g., memory hidden from an input/output  
25    bus or memory protected against DMA). In partial speculation mode, speculation is permitted, for example, for the benefit of correctly handling "architectural faults"

(exceptional conditions that do correspond to observable events in the abstract machine mentioned above).

In an embodiment in which microprocessor 103 is a microprocessor that  
5 utilizes a translator module and an interpreter module as described above,  
interpretation of non-native instructions can be performed in full or partial  
speculation modes, while translations are executed in full speculation mode.

In one embodiment, a speculation boundary corresponds to a commit  
10 point. An example of microcode according to such an embodiment is illustrated  
by the following:

	load_ss( );	[1. load stack segment]
	load_cs( );	[2. load code segment]
15	cmit ( );	[3. speculation boundary]
	load_ds( );	[4. load data segment]
	cmit ( ).	[5. speculation boundary]

In the example above, it is assumed that the ISA definition is such that the  
20 final state after a failure to load the code segment at line 2 would contain the  
stack segment as it was prior to entering this microcode sequence, whereas after  
a failure to load the data segment on line 4 the final state would contain the stack  
and code segments loaded at lines 1 and 2. In this example, if the speculation  
system cannot be used, additional tests would need to be added to verify that the  
25 code segment load on line 2 would complete successfully before the stack  
segment load on line 1 is attempted.

Thus, in response to an event, partial speculation mode permits some speculative operations to be performed, in lieu of suspending all speculative operations in response to the event. Note that in some situations, all speculative operations may still be suspended. For example, should another event occur  
5 during operation in partial speculation mode, all speculative operations may be suspended. Thus, according to the embodiments of the present invention, at least three speculation modes are permitted: full speculation, partial speculation, and all speculation suspended.

10 It is appreciated that other types of full and partial speculation modes can be defined. That is, for example, a full speculation mode or a partial speculation mode can be defined that includes operations other than those mentioned herein. Similarly, different levels of partial speculation modes can be defined. For example, the highest level of partial speculation could permit operations that are  
15 a subset of the operations permitted with full speculation. The next level of partial speculation could be a different subset of the operations permitted with full speculation. Alternatively, the next level of partial speculation could be a smaller subset of the larger subset of operations associated with the highest level of partial speculation.

20

In step 5 of Figure 3, execution in partial speculation mode continues until the event is handled or until some other condition is satisfied. For example, in step 1, a count can be made of the number of instructions executed in full speculation mode. When the same number of instructions are executed in step 4  
25 in partial speculation mode, then a return to full speculation mode can be made.

Speculation boundary 311 corresponds to a memory state that exists during the execution of step 5. In one embodiment, speculation boundary 311 corresponds to another commit point, subsequent to the commit point associated with speculation boundary 310.

5

Figure 4 is a flowchart 400 of a method providing partial speculative operation according to an embodiment of the present invention. Figure 5 is a flowchart 500 of another method providing partial speculative operation according to an embodiment of the present invention. Although specific steps are disclosed in flowcharts 400 and 500, such steps are exemplary. That is, 10 embodiments of the present invention are well suited to performing various other steps or variations of the steps recited in flowcharts 400 and 500. It is appreciated that the steps in flowcharts 400 and 500 may be performed in an order different than presented, and that not all of the steps in flowcharts 400 and 15 500 may be performed.

Referring first to Figure 4, in step 410, operations are executed in full speculation mode. Operations are executed forward from a speculation boundary (e.g., speculation boundary 310 of Figure 3). The speculation 20 boundary represents a state of the microprocessor (e.g., microprocessor 103 of Figures 1 and 2).

In step 420 of Figure 4, an event (e.g., a fault, an I/O read, or a DMA request) is detected.

25

In step 430, the initial microprocessor state is restored by rolling back to the speculation boundary (e.g., speculation boundary 310 of Figure 3).

In step 440 of Figure 4, operations are executed from the speculation boundary in partial speculation mode. Once the event is handled, or once some other type of condition is satisfied, execution can return to full speculation mode.

5 Alternatively, speculation may be suspended in entirety should another event occur during partial speculation mode.

Referring now to Figure 5, in step 510, operations are executed while in a first mode of speculative operation (e.g., a full speculation mode) that permits a  
10 first set of speculative operations.

In step 520, an event is experienced. The event is of a type such that execution in the first mode of speculative operation is no longer permissible. Full speculation mode may interfere with the externally visible behavior of  
15 microprocessor 103 (Figures 1 and 2), and so full speculation mode is suppressed for some events such as DMA requests, I/O reads and faults.

In step 530 of Figure 5, in response to the event, the first mode of operation is exited and a second mode of operation is entered. In the second  
20 mode (e.g., a partial speculation mode), a second set of speculative operations is permitted, the second set being a subset of the first set. The second set of speculative operations include those operations that are limited to memory that is private to the microprocessor (e.g., memory hidden from input/output buses or memory protected against DMA), such as CPU register 102 of Figures 1 and 2.

25

In step 540 of Figure 5, after the event is handled (or after some other condition is satisfied), execution can return to the first mode of operation (e.g., to full speculation mode).

5 In summary, according to the various embodiments of the present invention, a partial speculation mode of operation is introduced. The partial speculation mode is in addition to a full speculation mode and a mode in which speculation is suspended in entirety. While in partial speculation mode, processor register operations can be speculated. Furthermore, in partial  
10 speculation mode, memory operations under exclusive control and use of a microprocessor can be speculated. As such, microcode can be written that makes use of the speculation system.

Embodiments of the present invention have been described. The  
15 foregoing descriptions of specific embodiments of the present invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and obviously many modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best  
20 explain the principles of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the Claims appended hereto and their equivalents.